

AMENDMENTS TO SPECIFICATION

On page ~~7~~<sup>6</sup>, please replace lines 14-17 with the following:

- A<sup>1</sup>
- RTT time 117, 118 from the latency probe to LDNS.
  - ASN (Autonomous System Number) routing information derived from Border Gateway Protocol (BGP).
  - Dynamic hop count 117, 118 from the latency probe to LDNS.

On page 8, please replace lines 1-5 with the following amended paragraph:

A<sup>2</sup>

The latency probe 104, 106 uses a UDP Reverse Name Lookup and Traceroute to determine RTT and dynamic hop count 117, 118. Reverse Name Lookup is a standard DNS query that specifies a client IP address and asks for the client name. Traceroute is a specific format of a packet that is sent between routers to indicate if a packet has reached a destination.

On page 9, please replace lines 5-10 with the following amended paragraph:

A<sup>3</sup>

The address of the client 208 is masked to the IP prefix of the ASN 203. For example, if the address of the client 208 is 4.10.20.30 then the address is masked to the ASN 203 which is 4.0.0.0. The mask can vary depending on the higher level server granularity desired. For example, the first eight bits may be masked to the higher level server DNS 207, i.e., the address would then be masked to 4.10.20.00. This can be adjusted depending on the size of the network.

On page 9, please replace lines 22-27 with the following amended paragraph:

A<sup>4</sup> Referring to Fig. 3, a latency management table is used by the invention. The table contains the following fields: IP address 301; BGP hop count 302; and Trace data 303. The IP address field 301 contains the IP addresses that the server is responsible for. BGP hop counts 302 is taken from the BGP routing table for the particular IP address. The Trace data field 303 contains the latency and dynamic hop count information obtained by the latency probes.

Please replace page 9, line 29 - page 10, line 5 with the following amended paragraph:

A<sup>5</sup> The BGP hop count 302 is used by the SPDNS when the first request from an LDNS comes in for a particular IP address. No previous connection has been established at this point. This is because the dynamic hop count to the BG takes some time to actually measure. Subsequent requests use the actual dynamic BG hop count measured by the system located in the Trace field 303. The latency measurement is also taken and the combination of the dynamic hop count and latency measurement in the Trace field 303 is used to determine which SPDNS is closest to the client.

On page 10, please replace lines 7-12 with the following amended paragraph:

A<sup>6</sup> The invention does not need to get absolute hop counts from the various probe servers to the LDNS. It is sufficient to determine the relative (dynamic) hop counts between the probe servers and the LDNS and use this information in arriving at relative latency metrics between the client and the probe servers. This information can then be used by the Speedera DNS to choose between the service points and locate a server that is closest to the client.

Please replace page 10, line 26 - page 11, line 6 with the following amended paragraph:

A<sup>7</sup> Referring to Fig. 5, the two probe servers (Probe1 502 and Probe2 503) have initiated probes to the LDNS 501. The LDNS 501 returns a response to each of the probe servers. Here, the LDNS 501 sets the TTL in the response IP packet to  $R$  504, 505. The TTL will be decremented by 1 each time the packet passes through a router (hop). In this example, the packets go through  $H1$  hops 506 between LDNS 501 and Probe1 502 and  $H2$  hops 507 between LDNS 501 and Probe2 503. The TTL of the response packet that arrives at Probe1 502 will be  $(R-H1)$  and which arrives at Probe2 503 will be  $(R-H2)$ . The TTL of the response at Probe1 502 and Probe2 503 will be in inverse relation to the number of hops 508, 509 (The fewer the number of hops, the higher the TTL). Thus, a relative dynamic hop count can be arrived at by subtracting the TTL in the response packet from a fixed value  $R$ .

On page 11, please replace lines 8-11 with the following amended paragraph:

A<sup>2</sup> The latency metric is a weighted combination of the RTT and the dynamic hop count (e.g.,  $(RTT * w1) + (dynamic\ hop\ count * w2)$ , where  $w1$  and  $w2$  are relative weights). The latency metric is used to determine the server that is the most efficient for accessing a client. The invention precisely determines the hop count metric between client and server.

Please replace page 11, line 30 - page 12, line 6 with the following amended paragraph:

A<sup>9</sup> Request for latency metrics activate the Send Latency Metric module 601 to lookup the latency metric for the requesting LDNS's client. The Send Latency Probe module 603 sends latency probes to the IP addresses in the Latency Management Table 604. The IP addresses of clients are masked so the latency probes are sent to higher level servers as detailed above. Packets sent in response to the latency probes sent by the Send Latency

X<sup>9</sup>  
A Probe module 603 are received by the Receive Response Packet module 602. Hop  
Dynamic hop count and latency data are stored into the Latency Management Table 604.

---

On page 12, please replace lines 8-15 with the following amended paragraph:

---

A<sup>10</sup>  
The Send Latency Metric module 601 uses the information in the Latency Management Table 604 to determine the latency metric from the resident POP to the requesting LDNS's client before sending the latency metric to the requesting server. The Send Latency Metric module 601 uses the BGP hop count in the Latency Management Table 604 for its calculations upon the first request for an IP address. The latency metric is calculated for subsequent requests of IP addresses by the Send Latency Metric module 601 using the dynamic hop count and RTT data obtained from the Receive Response Packet module 602.

---

On page 12, please replace lines 17-23 with the following amended paragraph:

---

A<sup>11</sup>  
Latency metrics from POPs are received by the Receive Latency Metrics module 607. The latency metrics are sent to the Determine Optimal Server module 608. The Determine Optimal Server module 608 gathers the expected latency metrics and uses the inverse relationship of the dynamic hop counts in a weighted combination with the RTT to determine which latency metric indicates the optimal POP. The Determine Optimal Server module 608 then sends the address of the optimal POP to the requesting LDNS.

---